APSIM

An Asset Price and Portfolio Choice Simulator

William F. Sharpe

The APSIM Program and APSIM Cases

This document provides information on the APSIM simulation program. The executable version of the program, APSIM.EXE must be installed in a directory named c:\apsim on the user's windows computer. Each CASE is contained in a special Excel workbook that contains macros designed to interact with and execute the APSIM.EXE program. To prepare a new case, the user can save an old case workbook under a new name, make the desired changes, then press the PROCESS button on the case worksheet.

It is important to avoid copying a PROCESS button from one workbook to another. Each workbook's process button, when pressed, will take the inputs from that workbook, process them, then put the outputs on that workbook. Thus if the button from workbook a is copied to workbook b, pressing the button in workbook b will process workbook a. The simple rule is: never copy a PROCESS button from one workbook to another. Instead, save a workbook under a new name, then make any desired changes to the information on the case sheet of that workbook. It is, of course, perfectly safe to copy tables, etc. from one workbook to another or from one worksheet in a workbook to another in that workbook.

The simulator is used extensively in the author's book, *Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice,* Princeton University Press, forthcoming. Until its publication, a draft version will be available at www.wsharpe.com.

The apsim program, its source code and the cases used in the book can be downloaded from www.wsharpe.com.

The best way to use apsim is to begin with one of the cases from the book, then make modifications as desired. This document is designed to provide detailed documentation of the program's features, requirements and capabilities. It is best used as a reference document.

There are six main sections. Section 1 describes the possible inputs for an apsim case. Section 2 describes the procedures used for trading in the simulation. Section 3 describes the information showing the details of inputs actually used for the simulation. Section 4 covers the equilibrium output tables produced by the simulation and section 5 the associated graphs of the results. Section 6 describes the information on the log worksheet.

The table of contents is shown on the next two pages.

1. Inputs

The inputs for a simulation case are described on the "Case" worksheet of an APSIM workbook. To create an APSIM workbook, make a copy of one of the standard cases distributed with the system. The workbook includes the required Visual Basic (macro) commands and buttons to run the APSIM se imulator.

The case description consists of a series of tables with column names and row names. The name of a table must be in the upper left-hand corner and correspond to one of the reserved names shown below, including the final colon. There must be a blank row above and below each table. Unless the table starts in column A, there must be a blank column to its left. There must also be a blank column to its right. Tables must not start to the right of column Z. Tables may be included in any order. Colors and borders are not required for case inputs but are provided automatically for outputs.

The body of each table must be numeric – that is, every entry in the body of the table must be a numeric value or completely blank.

The full set of tables used for a simulation can be input directly. Alternatively, some shorter versions may be used. In either event, a detailed set of all tables to be used in the simulation is produced by the simulator's pre-processor and shown after the simulation is complete on the "Detail" worksheet of the case workbook (described in section 3, below)..

The simulator's preprocessor checks to insure that most of the conditions described in this document are met. If an error is found, the program will return an error message, show it on the "Log" worksheet, and display that worksheet. Any error will generate a message and halt the program, thus it may take several steps to remove multiple errors.

The simplest set of tables is included in the Case1.xls workbook, which uses only the tables that must be present. We describe these tables first.

1.1 The Securities Table

The securities table shows the securities to be used in the simulation and their payoffs in each of a series of possible states of the world. Each column represents a security and each row a state of the world. The first state must represent the present (now). The remaining rows must represent alternative states of the world at a single future date. The first security must pay 1 in the first state of the world and 0 otherwise; it represents current consumption. The second security must pay 1 in every future state of the world and 0 now; it represents a riskless security paying 1 unit at the future date with certainty.

All entries in the securities table must be positive or zero. Thus only limited liability securities can be included.

The names of the securities and the names of the states in the securities table will be used for all outputs.

In Case 1 the securities table is:

Securities:	Consume	Bond	MFC	HFC
Now	1	0	0	0
BadS	0	1	5	3
BadN	0	1	3	5
GoodS	0	1	8	4
GoodN	0	1	4	8

1.2 The Portfolios Table

The portfolios table shows the number of units of each security held by each of the investors. Each column represents a security, in the same order as in the securities table. Each row represents an investor. The names of the investors in this table will be used for all outputs.

In Case 1 the portfolios table is:

Portfolios:	Consume	Bond	MFC	HFC
Mario	49	0	10	0
Hue	49	0	0	10

The number of columns in the portfolios table must be the same as the number in the securities table.

1.3 The Probabilities Table

This probabilities table shows the actual probability of each state. The states must be listed in the same order as in the securities table, although in the probabilities table the states are in columns. The probability for the first state (Now) must equal 1.0. The probabilities for the remaining (future) states must sum to 1.0. Each probability must be greater than 0.0 and less than or equal to 1.0.

In Case 1 the probabilities table is:

Probabilities:	Now	BadS	BadN	GoodS	GoodN
Probability	1	0.15	0.25	0.25	0.35

The number of columns in the probabilities table must equal the number of rows in the securities table.

1.4 The Preferences Table

The final required table is the preferences table. Each row represents an investor, listed in the same order as in the portfolios table. For Case 1 there are only two columns. The first shows each investor's time preference, the second his or her risk aversion.

Time preference indicates the extent to which the investor discounts consumption in the future period relative to consumption in the present, other things equal. Thus a time preference of 0.96 indicates that $0.96 of certain future consumption would be as desirable for the investor in question as $1.00 of present consumption.

Risk aversion indicates an investor's attitude towards risk. This must be a positive number. The larger the number, the greater the investor's aversion to risk. A value between 2.0 and 3.0 is typical for an investor with average risk aversion.

More information on preferences is found in a later section of this document.

In Case 1 the preferences table is:

Preferences:	Time	RiskAversion
Mario	0.96	1.5
Hue	0.96	2.5

1.5 The Discounts Table

The four tables in Case 1 are sufficient to describe a case that can be processed by APSIM. However, much more detail can be included using additional tables. The following sections provide information on such inputs.

The preferences table used in base case 1 contains information about two aspects of each investor's preferences, dealing with time and risk. Each of these can be specified in more detail. The two tables that include all the information about an investor's preferences are the discounts and utilities tables. If a preferences table is utilized, these tables are generated and filled in automatically. Otherwise, both must be included instead of a preferences table.

A detailed discounts table corresponding to the information in the preferences table in case 1 is:

Discounts:	Now	BadS	BadN	GoodS	GoodN
Mario	1.00	0.96	0.96	0.96	0.96
Hue	1.00	0.96	0.96	0.96	0.96

The discounts table has one column for each state (in the same order as in the securities table) and one row for each investor (in the same orders as in the portfolios table). If desired, the discounts for an investor can differ in different future states, allowing for the representation of state-dependent preferences. The entries in the first column must all equal 1, since the present serves as numeraire, and all entries must be positive.

A short version of the discount table may be used if desired. In this case only one column is included, indicating for each investor the discount rate to be used for all future states. These values must all be positive. The simulator will then use this information to prepare the full discounts table to be used in the simulation.

1.6 The Utilities Table

The other detailed investor preference table provides information about each investor's utility function. A utilities table containing the information used in Case 1 is:

Utilities:	Type	Parameter
Mario	2.00	1.50
Hue	2.00	2.50

This indicates that for the entire range of possible consumption the first investor (Mario) has a utility of type 2 with a parameter of 1.5 and the second investor (Hue) has a utility of type 2 with a parameter of 2.5.

There are four possible types of utility function. Each investor must be given one of these types with the associated parameter values but investors may have different types of utilities. This may cause the table to have rows with different numbers of parameters, but blank entries may be used to the right of parameter values for rows with fewer entries.

In each case the marginal utility given by the formula utilized is multiplied by the discount factor for the state in question to give the "complete" marginal utility.

The types of utility are described in the following sections.

1.6.1 Type 1 Utility: Quadratic

Type one specifies a quadratic utility function, and hence a linear marginal utility function. Thus utility (u) is related to consumption (X) by:

$$u(X_s) = a + bX_s - cX_s^2$$

Marginal utility (m) is thus related to consumption by:

$$m(X_s) = b - 2cX_s$$

To avoid zero marginal utility a linear function is used over only part of the range of possible levels of consumption. Only one parameter is required, indicating a "near satiation" level of consumption. One investor's inputs for Case 3 provides an example:

Utilities:	Type	Parameter
Quentin	1	200

The marginal utility function is initially taken to be a linear function with a marginal utility of 100 for a consumption of 0 and a marginal utility of 1 for the near-satiation level of consumption (here, 200). For levels of consumption greater than or equal to the near-satiation level, an exponential utility function is utilized. This function is constructed to have the same marginal utility (1) and the same slope (dm/dx) at the near-satiation level as the linear marginal utility function. This insures that marginal utility will be positive at all levels of consumption.

Since consumption is never allowed to fall to or below zero (due to minimum consumption requirements described later) there is no need to alter the linear marginal utility function used for levels of consumption below the input near-satiation level.

1.6.2 Type 2 Utility: Constant Relative Risk Aversion

Type 2 utility is used when only a preferences table is utilized. It can also be specified for one or more investors in a utilities table. Such a function displays Constant Relative Risk Aversion (CRRA). The logarithm of marginal utility (m) is taken to be a linear function of the logarithm of consumption (X). Thus:

$$\ln(m) = a - b\ln(X)$$

Since the value of a will not affect investors' decisions, it is taken to be zero. Thus only one parameter (b) is required to specify a specific CRRA investor's utility function.

The utilities for the investors in Case 1 could be input as follows:

Utilities:	Type	Parameter
Mario	2.00	1.50
Hue	2.00	2.50

Here, Mario's CRRA utility function has a b value of 1.50 and Hue's has a b value of 2.50. In a graph with the log of marginal utility on the vertical axis and the log of consumption on the horizontal axis, each investor's marginal utility will plot as a straight line, with Mario's line having a slope of -1.50 and Hue's a slope of -2.50.

The value of b is commonly referred to as the investor's relative risk aversion coefficient.

1.6.3 Type 3 Utility: Piecewise Constant Relative Risk Aversion

Type 3 utility functions are composed of segments, each of which displays constant relative risk aversion. In a graph with the log of marginal utility on the vertical axis and the log of consumption on the horizontal axis, each segment plots as a straight line with the segments connected. Each segment applies over a range of consumptions, with the segments listed in order of increasing consumption.

The following utilities table from Case 21 provides an illustration.

Utilities:	type	ra1	Cons1	ra2	Cons2	ra3
Karyn	3	3.00	98	50	98.98	3.00

Here, Karyn has a constant relative risk aversion of 3.00 for levels of consumption from 0 to 98. She has a constant relative risk aversion of 50 for levels of consumption from 98 to 98.98, then a constant relative risk aversion of 3.00 for levels of consumption from 98.98 to infinity.

To provide a scale, the marginal utility is set to equal 1 at the level of consumption at the end of the first segment (98 in this case). Given the requirement that each adjoining segment have the same marginal utility at the common level of consumption, this determines the equations for all the marginal utility functions.

1.6.4 Type 4 Utility: Hyperbolic Absolute Risk Aversion (HARA)

Type 4 utility exhibits hyperbolic absolute risk aversion. Such a HARA function can be written as:.

$$m = (X - M)^{-b}$$

Here, m is marginal utility and M is a minimum level of consumption at which the investor's marginal utility becomes infinite. The parameter b represents the investor's risk aversion vis-à-vis consumption in excess of the minimum level M.

Case 4 provides an example. One of the investor's preferences are:

Utilities:	Type	MinCons	RiskAversion
David	4	20.00	1.00

Here the minimum acceptable consumption (M) is 20, and the coefficient of risk aversion for consumption above this level (b) is 1.00.

As described below, apsim requires that each investor's consumption be above a required level in every state. For investors with utilities of type 1, 2 or 3, this minimum is set to a universal amount specified as a control. However, for investors with utilities of type 4, the required minimum is equal to the minimum consumption indicated in the utilities table plus the standard required minimum specified as a control. If the initial portfolio for an investor leads to consumption in any state below the required minimum, the simulation is not performed. Otherwise, the simulation proceeds, but no trades are allowed that will violate the required minimum levels of consumption.

1.7 The Collateral Table

The collateral table shows each investor's income in each of the states from non-portfolio assets that can serve as collateral for portfolio positions. The table must have as many rows as there are investors and as many columns as there are states. Since only future income from collateral is shown, the first column must contain zeros. Otherwise, only zero or positive entries are allowed.

The table from Case 9 provides an example.

Collateral:	Now	BadS	BadN	GoodS	GoodN
Mario	0	25	15	40	20
Hue	0	15	25	20	40

A short form is also allowed. In this case only one column is included. The entry in each row indicates the income from collateral in all future states for the investor in question. The preprocessor then provides a full table with the desired entries.

Collateral income can be utilized to guarantee any net obligations from portfolio positions. Thus an investor can take portfolio positions that lead to a net negative cash flow in a state as long as the income from collateral in that state is larger than the magnitude of the required negative cash flow. Put somewhat differently, the simulator enforces a credit check at each stage to prohibit transactions that would result in the sum of an investor's income from collateral and portfolio cash flow to be negative.

If every investor will receive the same income from collateral in every future state, a short form may be utilized, with one column showing the investors' future incomes from collateral.

If no collateral table is included in the case description, one is created by the preprocessor with all entries equal to zero.

1.8 The Salaries Table

This is a table showing each investor's income from salary in each of the states. It must have as many rows as there are investors and as many columns as there are states. Since only future salary income is shown, the first column must contain zeros. Otherwise, only zero or positive entries are allowed.

The table from Case 10 provides an example.

Salaries:	Now	BadS	BadN	GoodS	GoodN
Mario	0	30	15	45	20
Hue	0	15	25	20	40

A short form is also allowed. In this case only one column is included. The entry in each row indicates the salary in all future states for the investor in question. The preprocessor then provides a full table with the desired entries.

Salary income cannot be utilized as collateral, since an investor can declare bankruptcy in any state in which net obligations from collateral plus portfolio positions are negative, thereby keeping the entire salary income in that state.

If every investor will receive the same salary in every future state, a short form may be utilized, with one column showing the investors' future salaries.

If no salaries table is included in the case description, one is created by the preprocessor with all entries equal to zero.

1.9 The Predictions Table

APSIM allows investors to have different predictions, that is, estimates of the probabilities of alternative states of the world. Using information provided in the case description, APSIM generates a detailed prediction table showing the probabilities of the states used by each investor. If no predictions table is included in the case description (as, for example, in Case 1) each investor's predictions are set to equal to the actual probabilities of the states. For Case 1 the following predictions table is generated by the apsim preprocessor:

Predictions:	Now	BadS	BadN	GoodS	GoodN
Mario	1	0.15	0.25	0.25	0.35
Hue	1	0.15	0.25	0.25	0.35

If desired, however, a predictions table can be provided explicitly in the case description. It must have a row for each investor and a column for each state. The entries indicate each investor's probability assessments for the states. The entries must meet the same standards as the probabilities table. For each investor the probability for state 1 (now) must be 1.0, the sum of the probabilities for the future states must equal 1.0 and each probability must be greater than 0.0 and less than or equal to 1.0.

1.10 The Information Table

An alternative way to obtain diverse predictions assumes investors have different sources of information. The information table from Case 14 provides an example.

Information:	Prior Wt	Samples
Mario	0.01	100
Hue	0.01	100

In this case the actual probabilities used for future states for each investor will be a weighted average of two sets of probabilities. The first set of probabilities will be taken from the investor's row in the prediction table or, if no prediction table is given, from the actual probabilities shown in the probabilities table. The second set of probabilities will be based on a sample drawn from the actual probabilities.The entry in the second (Samples) column indicates the number of samples to be drawn to obtain the second distribution. Finally, the two distributions will be combined, with the prior distribution given the weight shown in the first (PriorWt) column and the sampled frequency distribution given a weight equal to 1 minus the weight on the prior distribution.

The weight for the prior distribution must be greater than 0.0 and the number of samples must be less than or equal to 1. This insures that the predictions generated for each investor will have a positive probability for every state.

If an information table is included in the case description a predictions table is generated by the apsim preprocessor using the procedure described above. The resulting predictions are then shown in the Detail worksheet (described below).

A set of random numbers is used to generate each sample. The random number generator uses a "seed" that can be set in the Controls table (described below). If the random seed in the controls table is set to a positive number the same predictions will be generated every time the case is processed. If the random seed is set to 0, every time the case is processed, a different set of predictions will be generated. If no random seed is set, the default value of 0 is utilized.

There is an optional long form for the information table as well. The following table provides an example.

Information:	Prior Wt	Prior dist	Samples	Sample Dist
Mario	0.01	1	100	2
Hue	0.02	2	200	1

In this version the distributions to be used for each investor are specified explicitly. There are two alternatives:

1. the actual probabilities shown in the probabilities table
2. the probabilities shown for the investor in the predictions table

The short form is equivalent to the long form with prior distribution 2 (the investor's predictions) and sample distribution 1 (the actual probabilities).

If a predictions table is not provided on the case worksheet, each investor's predictions are taken to equal the actual probabilities.

As with the short form, the weight on the prior distribution must be greater than 0 and less than or equal to 1. This insures that the probability for each state will be positive. The number of samples must be greater than or equal to 1.

1.11 The Control Table

The apsim program uses five control parameters. These are contained in a controls table, which is generated by the apsim preprocessor if no table is included on the Case worksheet. The default controls table is shown below.

Controls:	Parameter
Pct Price Spread	0.001
Complete Market	0.000
Random # Seed	0.000
Min. Consumption	1.000
Number of Prices	1.000

If a controls table is included on the Case worksheet it may have all five rows. Alternatively, a table with fewer rows may be provided and the preprocessor will complete the table using the default values for the remaining rows. In either case, rows must be given in the order shown above.

The controls are described below.

1.11.1 Percent Price Spread

This parameter is used to provide a threshold for actual trades. An investor will not buy a security unless his or her reservation price is greater than the trade price times (1+.01s), where s is the percent price spread. Similarly, an investor will not sell a security unless his or her reservation price less than the trade price times (1-.01s).

If the requirements for buying or selling a security are met, subject to minimum consumption constraints, at a trade price of p an investor will submit a bid or offer for the number of units of a security that would make his or her reservation price fall between p(1.+01s) and p(1-.01s).

This parameter provides an opportunity to trade simulation speed for accuracy. The greater the percent price spread the less accurate the results but the more quickly equilibrium will be achieved. In some cases it may be useful to run a case with a relatively large percent price spread to see the general nature of the equilibrium. If the results are interesting, the case can later be run with a lower percent price spread to obtain more accurate results.

The percent price spread can also be considered a crude surrogate for costs and/or frictions in the market-making process. However, aggregate consumption in each state is the same as in the initial conditions, so there will be no reflection of dead-weight costs associated with trades.

The percent price spread must be positive and less than or equal to 100.

1.11.2 Complete Markets

This parameter specifies whether or not a market in state claims is to be opened after trading is finished in the market for securities. The value must be either 0 or 1. A value of 0 indicates that trading in state claims will not take place, while a value of 1 indicates that such trading will take place. The default value is 0 (no complete markets).

1.11.3 Random Number Seed

The random number seed is utilized to generate random numbers if required to generate predictions based on an information table (if no information table is present, this parameter is not utilized).

The random number seed must be zero or positive. An integer number is preferred but not required. Different seeds will generate different predictions. If the seed is positive, the same predictions will be generated every time the case is processed. On the other hand, if the seed is 0, different predictions will be generated each time the case is processed.

The default random number seed is 0 so that different predictions will be generated each time a case with an information table is processed.

1.11.4 Minimum Consumption

For investors with utility of type 1, 2 or 3, the minimum consumption specifies the smallest amount that the investor is allowed to consume in any state. For an investor with utility of type 4 the minimum consumption allowed in any state is equal to the minimum specified for that investor in the utilities table plus the minimum consumption specified in the control table. If investor preferences are specified in a preferences table (as in Case 1) utility of type 2 is utilized so in every state each investor must have more than the minimum consumption specified in the control table.

If any investor's initial position provides consumption in any state that is below his or her required minimum consumption, the apsim program will not process the case. When a case is processed, no trade is allowed that would cause an investor to violate the minimum consumption in any state.

1.11.5 Number of Prices

The number of prices indicates the number of iterations used by the market maker before setting the price at which investors submit bids and offers for actual trades. In all the cases used in the author's book this is set at 1. This specifies that the market maker should use the simple price discovery procedure described in the next section, setting the trade price equal to the average of two values: (1) the mean of the reservation prices for those able to purchase the security without violating minimum consumption requirements and (2) those able to sell the security without violating minimum consumption requirements. If the number of prices parameter is set to a greater value, the more complex procedure price discovery procedure described in the next section is utilized.

The number of prices must be a positive integer. The default value is 1.

2. Trading

Apsim uses a stylized version of markets made by market makers. The program simulates a number of *rounds* of trading. Each round consists of an attempt to make a market in each of the future securities (all but the first security, which represents current consumption). There are two phases for each security. The first is price discovery. The market maker poses questions to the investors in an attempt to find a price at which there will be both demand and supply for trading the security. If this phase is successful, a market is conducted for the security at the price selected in the initial phase. Investors submit bids for quantities demanded or supplied. The market maker then finds the smaller of the quantity demanded and the quantity supplied, clearing the market by assigning each investor all or a proportion of the shares in his or her bid or offer. If demand exceeds supply, all sellers are able to sell the desired shares and each buyer receives a proportion of his or her submitted bid, with the proportion equal to the ratio of total supply to total demand. If supply exceeds demand, buyers are able to buy their desired shares and each seller is able to sell a proportion of his or her submitted offer, with the proportion equal to the ratio of total demand to total supply.

When submitting a bid or offer for units of a security an investor will choose a number of units of a security that would make his or her reservation price fall between p(1.+01s) and p(1-.01s), where p is the trade price and s is the percent price spread control.

When a round of trades has been completed, the market maker tallies the number of units of securities traded in the round. If the sum equals zero, the trading process is stopped since equilibrium in the securities market has been attained. If, however, some trades were made in the round, another round of trades is conducted. The process continues until equilibrium is reached.

There are two possible procedures for the price discovery phase in which the market maker determines the trade price. In most cases the simple price discovery process will suffice but in some cases the more complex approach may be desirable. Each is described below,

2.1 Simple Price Discovery

In the simple price discovery procedure, each individual who is able to sell the security in question submits his or her reservation price as a minimum sales price – above that price the investor will sell at least some shares. Each individual who is able to buy the security submits his or her reservation price as a maximum bid price – below that price the investor will buy at least some shares. If an investor would violate a credit constraint by selling any shares of a security, he or she does not submit a sales price. Similarly, if an investor would violate a credit constraint by buying any shares of the security, he or she does not submit a bid price.

If the largest maximum bid price for a security exceeds the smallest minimum ask price by the maximum percentage price control the market maker sets a trade price halfway between (1) the mean of the bid prices within the range from maximum bid to minimum ask (inclusive) and (2) the mean of the ask prices within the range from maximum bid to minimum ask (inclusive).

Simple price discovery is used for all the Cases in the author's book.

2.2 Complex Price Discovery

In some cases it may be desirable to utilize a more complex price discovery process in which multiple possible trade prices are examined, with investors asked to indicate the sizes of their bids or offers at each price. The market maker then selects from these prices the one that will produce the largest number of units actually traded. This is then announced as the trade price for the market.

This procedure is invoked if the number of prices control is greater than 1. If, for example, the number of prices is set to 5, five alternative prices will be considered. The first is the one that would be used for a simple price discovery. Subsequent prices are set halfway between a high price (h) and a low price (l). In the second iteration these prices are the highest bid and lowest ask price, respectively. In each iteration the number of units that could be traded is determined. Then either the high or low price is changed. If there would be excess demand at this price, it replaces the previous low price (l). If, on the other hand, there would be excess supply at this price, it replaces the previous high price (h). The process is repeated until the specified number of prices is obtained or the two prices (h and l) are equal. Of the prices analyzed, the one for which the number of units traded would be the greatest is selected as the actual trade price.

Complex price discovery requires more information from investors and more time for price discovery than simple price discovery. However, it may lead to fewer rounds of actual trading. Since the procedures used for trading affect the final distribution of wealth, different approaches to price discovery may lead to different equilibrium conditions. In general complex price discovery may provide more realistic results when there are large differences between the quantities demanded and supplied at the prices generated using the simple price discovery procedure. This is likely to be the case when there are great disparities in investors' initial endowments.

In actuality, market makers have some information about some investors' demand and supply curves but the information is partial at best. The trading and price discovery procedures available in apsim are intended to provide stylized versions of the actual processes used in security markets that should be sufficient to analyze the relationships that would hold once equilibrium is reached.

2.3 Complete Markets

In all cases trading is conducted in the securities shown in the securities table until equilibrium is reached. At this point the range of security reservation prices will be within the required threshold for all investors who are willing and able to buy and/or sell. However, this may or may not be true for the investors' reservation prices for state claims, since they have not been traded explicitly.

If the complete market control is set to 1 a second market is conducted after the market for securities reaches equilibrium. Then trading is conducted in state claims (and only in state claims) until a new equilibrium is reached. Rounds of trading are conducted until there are no trades in the most recent final round. For each state claim the trading process is exactly the same as that utilized for securities. The type of price discovery used for securities is utilized for state claims. The trading process is also the same.

To insure that there is equilibrium in both the securities market and the state claims market, after equilibrium is reached in the state claims market the simulation program conducts rounds of trading in both markets until no further trades can be made. In most cases the portfolio holdings will be the same as in a comparable incomplete market case, however, due to trading thresholds there may be slight differences.

When state claims are traded the equilibrium worksheet will include a table showing each investor's positions in state claims; otherwise this table is not included.

The default complete market control setting is 0, hence there is no trading in state claims unless the control is explicitly set to indicate that such trading is desired.

3. Detailed Inputs

After an apsim case has been processed, the tables actually used by the simulator are placed on the worksheet named Detail in the case workbook. Some of these will be the same as the corresponding table on the Case worksheet. Others may be default versions if no corresponding table was provided. Yet others will be expansions of short versions used in the Case worksheet or will be produced by preprocessing information from the Case worksheet. Any table from the Detail worksheet may be copied to the Case worksheet and used as is or modified as desired to create a new case.

While tables in the Detail worksheet are formatted to show a standard number of decimal places, the actual values are typically more precise. Standard Excel formatting instructions can be used to change the displayed precision for any entry.

Each table included on the Detail worksheet is described briefly below.

3.1 Securities Table

This will be the same as the securities table on the Case worksheet.

3.2 Portfolios Table

This will be the same as the portfolios table on the Case worksheet

3.3 Salaries Table

If a complete salaries table is included on the Case worksheet, it is repeated here. If a short form salaries table with only one column is included on the Case worksheet, this version shows the complete table. If no salaries table is included on the Case worksheet, all entries in this version will be 0.

3.4 Collateral Table

If a complete collateral table is included on the Case worksheet, it is repeated here. If a short form collateral table with only one column is included on the Case worksheet, this version shows the complete table. If no collateral table is included on the Case worksheet, all entries in this version will be 0.

l

3.5 Probabilities Table

This will be the same as the probabilities table on the Case worksheet

3.6 Predictions Table

If no predictions table is included on the Case worksheet, each row in this table will equal the probabilities in the probabilities table. If a predictions table is included on the Case worksheet and no information table is present, this will equal the predictions table on the Case worksheet. If an information table is included on the Case worksheet, the apsim preprocessor will create a set of predictions using the information provided using the procedure described in the earlier section on the information table; the resulting predictions will be shown in the predictions table on the Detail worksheet.

3.7 Discounts Table

If a complete discounts table is included on the Case worksheet, it will be repeated here. If a short form discounts table with only one column is included on the Case worksheet, this version shows the complete table. If instead a preferences table is included on the Case worksheet the time preferences in that table will be used to create the complete discounts table shown here.

3.8 Utilities Table

If a utilities table is included on the Case worksheet, it will be repeated here. If instead a preferences table is included on the Case worksheet the risk aversion coefficients in that table will be used to create the complete utilities table shown here.

3.9 Controls Table

If a complete controls table is included on the Case worksheet, it will be repeated here. If a partial controls table is included on the Case worksheet, the required additional rows with default values will be shown here. If there is no controls table on the Case worksheet the complete table with all default values will be shown here.

4. Equilibrium Outputs

After the apsim program has completed a simulation, information on the characteristics of the equilibrium is placed on the Equilibrium worksheet in the case workbook. If the case includes a market for state claims (specified by setting the complete markets control to 1) all the tables described below will be present. If there is no market for state claims (the default case) all the tables shown below with the exception of the Claims Table will be included.

While each table in the Equilibrium worksheet is formatted to show a standard number of decimal places, the actual values are typically more precise. Standard Excel formatting instructions can be used to change the displayed precision for any entry.

Each table included on the Equilibrium worksheet is described briefly below.

4.1 Portfolios Table

This table shows the final portfolio holdings after equilibrium has been attained. The top (Market) row shows the total holdings of each security; the remaining rows show the final portfolio holdings of each investor.

4.2 Claims Table (complete markets only)

This table is present only in complete market cases in which a market for state claims has been conducted after the market for securities has reached equilibrium. The entries show the final holdings of state claims. This table shows the final portfolio holdings after equilibrium has been attained in the market for such claims. The top (Market) row shows the total holdings of each state claim; the remaining rows show the final portfolio holdings of each investor. Since there are no state claims in the initial endowments, the market holdings of each claim will equal zero.

A complete market equilibrium with small holdings of state claims indicates that the original securities market allowed investors to allocate consumption relatively efficiently.

4.3 Consumptions Table

This table shows the amounts consumed in each state after equilibrium has been attained. All sources of consumption (portfolios, salaries and collateral) are included. The top (Market) row shows the total consumption; the remaining rows show the amounts consumed by each investor.

4.4 Security Prices Table

The security price table shows each investor's reservation price for each of the securities. Credit constraints are not taken into account when computing these prices. Thus even if an investor is unable to buy and/or sell any units of a security due to a minimum consumption constraint his or her reservation price will be calculated on the assumption that a purchase or sale is possible.

The top (Market) row in the security prices table shows the security prices set by the market maker in the last round of trading (in which no transactions were made). If the (default) simple price discovery process is being used this is equal to the average of the mean of the reservation prices of investors able to purchase shares and the mean of the reservation prices of investors able to sell shares. If the complex price discovery process is being used, this is the price established using the procedures described in section 2.

If minimum consumption constraints are not binding in the equilibrium, the investors' reservation prices for securities will differ by amounts determined by the percent price spread control and the market price will be close to an average of the investors' reservation prices.

4.5 State Prices Table

The state prices table shows each investor's reservation price for each of the state claims, whether or not trading in such claims has taken place. Credit constraints are not taken into account when computing these prices. Thus even if an investor would be unable to buy and/or sell any units of a state claim due to a minimum consumption constraint his or her reservation price will be calculated on the assumption that a purchase or sale is possible.

In a case involving an incomplete market, the Market state prices shown in the top row are determined by opening a market for each state claim but not completing any trades. In such a case the Market row shows the state prices set by the market maker for such a round of trading. If the (default) simple price discovery process is being used this is equal to the average of the mean of the reservation prices of investors able to purchase state claims and the mean of the reservation prices of investors able to sell state claims. If the complex price discovery process is being used, this is the price established using the procedures described in section 2.

In an incomplete market, investor's reservation prices for a state claim may differ. If such differences are small, the original securities market allowed investors to allocate consumption relatively efficiently.

In a case involving a complete market, the top row shows the Market state prices in the last round of trading. If minimum consumption constraints are not binding in the equilibrium, the investors' reservation prices for state claims will differ by amounts determined by the percent price spread control and the market price will be close to an average of the investors' reservation prices.

4.6 PPCs table

Each entry in this table is computed by dividing the corresponding entry in the state prices table by the actual probability for the state in question shown in the probabilities table for the case. For each state the table shows each investor's reservation price per chance and the market price per chance.

The entries in the top row are taken as the pricing kernel for the equilibrium.

4.7 Portfolio Values Table

This table shows the market value of each investor's security holdings as well as the total market values for each security. Each entry is computed by multiplying the corresponding entry in the portfolios table by the entry in the top row of the security prices table.

4.8 Portfolio Returns Table

This table shows the returns for each portfolio in each of the future states. For each investor the total value of all holdings of securities (other than the first security) is determined using the information in the portfolios values table. Then the consumption provided by the investor's portfolio in each future state is divided by the current value of that portfolio to determine the total return of the portfolio in that state. Consumption provided by salary and collateral is not included in the calculations.

The same procedure is used to compute the return on the Market, shown in the top row. The consumption provided by the market portfolio in each state is divided by the total value of the market holdings of standard securities (all but the first) to compute the market return in the state.

4.9 Security Returns Table

This table shows the returns for each security and for the market as a whole in each of the future states. For each security and state the payoff amount shown in the securities table is divided by the market security price shown in the top row of the security prices table to determine the total return for the security in that state. The top row repeats the information from the top row of the portfolio returns table.

4.10 Security Characteristics Table

The security characteristics table provides a number of statistics for each of the standard securities (all but the first) and for the overall market portfolio. With the exception of the last two (power beta and power alpha) all are widely used in the investment industry. The following sections provide brief descriptions.

4.10.1 Expected Return

Each entry in this column is computed taking a probability-weighted average of the corresponding returns in the future states shown in the security returns table, using the probabilities shown in the probabilities table. The result is the expected total return for the security, based on actual probabilities.

4.10.2 Expected Excess Return

Each entry in this column is computed by subtracting the expected (riskless) return on the second security (bond) from the expected return shown in the first column. The result is the expected excess return over and above the riskless rate for the security.

4.10.3 Standard Deviation of Return

Each entry in this column is computed using the returns in the future states shown in the security returns table for the security or market portfolio and the actual probabilities shown in the probabilities table. The standard deviation is the square root of the probability-weighted squared deviations of return from the expected return.

4.10.4 Sharpe Ratio

This is the ratio of the security or market's expected excess return to its standard deviation. If the standard deviation is zero, the Sharpe Ratio is shown as 0.

4.10.5 Beta

This is the ratio of the security or market's covariance with the market return divided by the variance of the market return. Equivalently, it is the slope coefficient in a regression in which the market return is the independent variable and the security or market return is the dependent variable. The covariances and the market variance are computed using the actual probabilities from the probabilities table.

4.10.6 Alpha

This is the difference between the security or market's expected excess return and its beta times the market's expected excess return. It provides a measure of the extent to which a security's expected excess return differs from that of the combination of the market portfolio and the riskless asset with the same beta as the security.

4.10.7 Power Beta

This is the ratio of (1) the covariance of the security or market's return with the total return of the market raised to the power shown as the risk aversion coefficient in the kernel approximation table (described below) divided by (2) the variance of the market's return with the return of the market to that power. It provides a measure of market-related risk based on the constant relative risk aversion pricing kernel function shown in the kernel approximation table.

4.10.8 Power Alpha

This is the difference between the security or market's expected excess return and its power beta times the market's expected excess return. It provides a measure of the extent to which a security's expected excess return differs from that of a combination of the market portfolio and the riskless asset with the same power beta as the security.

4.11 Portfolio Characteristics Table

This table provides the same statistics shown in the securities table for the portfolios held by investors and for the market portfolio. The calculation of each statistic is described in the appropriate section for the securities table.

4.12 Kernel Approximation Table

This table provides statistics obtained from a simple regression in which the independent variable is the logarithm of total consumption (shown in the top row of the consumptions table) and the dependent variable is the logarithm of PPC (shown in the top row of the PPCs table). Only the future states (all but the first) are used for the regression.

The first statistic (risk aversion) is the slope coefficient from the regression. The second is the exponential of the intercept from the regression. The third statistic is the r-squared value, which provides a measure of the extent to which the log-linear relationship fit the data points.

The kernel approximation table shows the parameters of the constant relative risk aversion marginal utility function that can best serve a proxy for the preferences of a "representative investor". The higher the value of r-squared, the more useful will be the associated approximation.

4.13 PPCs & Consumptions Table

This table repeats the information in the top (Market) rows of the consumptions table and the PPCs table. It is designed to show the relationship between the pricing kernel (PPC values) and total consumption.

4.14 PPCs & Market Returns Table

This table repeats the information in the top (Market) rows of the market returns table and the PPCs table. It is designed to show the relationship between the pricing kernel (PPC values) and market returns.

4.15 Sorted Portfolio Returns Table

This table contains the same information shown in the portfolio returns table but the columns are arranged in order of increasing Market returns (shown in the top row). It is designed to show the ways in which investor returns differ across states in which market returns increase.

5. Graphs

After the apsim program has completed its calculations the visual basic program in the case workbook posts the results to the Detail and Equilibrium worksheets, then creates six graphs of key results. Since different computers may have different graphic resolutions it is best to run each case on the computer on which it is to be analyzed in order to maximize the usefulness of the graphs.

It is possible to find the values for any point plotted in one of these graphs by placing the cursor over the point. It is also possible to find the locations on the Equilibrium sheet that provide the data for a series by clicking on the right button on the mouse within the chart area, then selecting Source Data.

Each of the graphs produced by apsim is described briefly below.

5.1 Returns

This graph plots each investor's returns on the vertical axis and the market portfolio's returns on the horizontal axis. The market portfolio's returns are also plotted. The data for the series are taken from the sorted portfolio returns table. Series 1 is the first investor listed, series 2 the second and so on. The final series shows the data for the market portfolio.

5.2 Capital Market Line

This graph plots the expected returns on the vertical axis and the standard deviations on the horizontal axis for all the securities and portfolios, including the market portfolio. The line is drawn through the points for the riskless security and the market portfolio. Points on the line can be obtained by combining the riskless security and the market portfolio; this is generally termed the Capital Market Line. Securities or portfolios that plot on the line have the same Sharpe Ratio as the market portfolio; those that plot above it have a higher Sharpe Ratio while those that plot below it have a lower one. The data are taken from the portfolio characteristics and security characteristics tables.

5.3 Security Market Line

This graph plots the expected returns on the vertical axis and the (standard) beta values on the horizontal axis for all the securities and portfolios, including the market portfolio. The line is drawn through the points for the riskless security and the market portfolio. Points on the line can be obtained by combining the riskless security and the market portfolio; this is generally termed the Security Market Line. Securities or portfolios that plot on the line have zero (standard) alpha values; those that plot above it have positive alpha values while those that plot below it have lower ones. The data are taken from the portfolio characteristics and security characteristics tables.

5.4 Pricing Kernel & Consumption

This graph plots the PPC values on the vertical axis and the amounts of total consumption on the horizontal axis. Each point represents a state, although two or more states may plot at the same point. The data are taken from the PPCs & Consumptions table. The points are connected with straight lines in the order appearing in the table strictly as a visual aid.

5.5 Pricing Kernal & Market Return

This graph plots the PPC values on the vertical axis and the market returns on the horizontal axis. Each point represents a state, although two or more states may plot at the same point. The data are taken from the PPCs & Market Returns table. The points are connected with straight lines in the order appearing in the table strictly as a visual aid.

5.6 Power Security Market Line

This graph plots the expected returns on the vertical axis and the power beta values on the horizontal axis for all the securities and portfolios, including the market portfolio. The line is drawn through the points for the riskless security and the market portfolio. Points on the line can be obtained by combining the riskless security and the market portfolio; this is termed the Power Security Market Line. Securities or portfolios that plot on the line have zero (standard) power alpha values; those that plot above it have positive power alpha values while those that plot below it have lower ones. The data are taken from the portfolio characteristics and security characteristics tables.

7. The Log

After the apsim program has finished executing it writes a message on the Log worksheet in the case workbook. If there are no errors in the case input this will simply indicate the time required for the simulation, for example:

APSIM 2.00: No errors. Processing time: 0.005 minutes

Processing time will differ from computer to computer and to a lesser extent from run to run on the same computer. The time shown on the log worksheet refers only to the time taken by the apsim program itself and does not include the time required to execute the visual basic instructions in the workbook.

If the inputs for a case are incomplete or violate some of the requirements indicated in this document the log file will contain a message indicating the first such problem encountered and the workbook will be positioned to show the message. If there are multiple problems with a case input file several attempts may be required to obtain a successful simulation.

The apsim program has been tested extensively to insure that it cannot interfere with other programs. To be doubly sure, internal checks are performed during processing. In the highly unlikely event that one of these checks indicates that processing should be halted to avoid any potential problem, processing will be stopped and the log file will show a MEMCHECK message. In such an event an email should be sent to that author with the Case worksheet attached so the problem can be identified.